

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (currently amended) A file system, comprising:

a plurality of servers configured to store file data as chunks; and

a master connected to the servers and configured to:

store namespace data that includes file identifiers for files for which the file data is stored as chunks,

store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond,

store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and

store location data that identifies which of the servers stores which of the chunks,

where the master is configured to:

communicate with the servers [[upon]] at startup of the master to identify the chunks stored by the servers, and

record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data.
2. (canceled)

3. (currently amended) The system of claim 1, ~~wherein~~ where the master is further configured to control placement of new chunks at the servers.

4. (currently amended) The system of claim 3, ~~wherein~~ where when controlling the placement of new chunks, the master is configured to:

identify one or more of the servers to store the new chunks based on at least one of utilization of the servers, prior chunk distribution involving the servers, network topology, or failure correlation properties associated with the servers, and
place the new chunks at the identified one or more servers.

5. (currently amended) The system of claim 1, ~~wherein~~ where the master is further configured to control redistribution of the chunks stored by the servers.

6. (currently amended) The system of claim 5, ~~wherein~~ where when controlling redistribution of the chunks, the master is configured to:

select a chunk to redistribute based on a current distribution of the chunks,
identify one or more of the servers to which to move the selected chunk, and
move the selected chunk to the identified one or more servers.

7. (currently amended) The system of claim 1, ~~wherein~~ where the master is further configured to monitor a state of the servers.

8. (currently amended) The system of claim 7, ~~wherein~~ where the master is configured to exchange heartbeat signals with the servers to determine the state of the servers.

9. (currently amended) The system of claim 8, ~~wherein~~ where the heartbeat signals include space utilization information.

10. (currently amended) The system of claim 7, ~~wherein~~ where the state of the servers includes information regarding the chunks stored by the servers.

11. (currently amended) The system of claim 10, ~~wherein~~ where the information includes version numbers of the chunks.

12. (canceled)

13. (previously presented) A master in a file system that includes the master connected to a plurality of servers, the master comprising:

means for communicating with the servers to identify file data stored by the servers as chunks;

means for storing, in a non-persistent manner, location information that identifies ones of the servers that store the chunks;

means for updating the location information by periodically instructing the servers to identify the data stored by the servers;

means for storing namespace data that includes file identifiers for files for which the file data is stored as chunks by the servers;

means for storing mapping data that maps the file identifiers to the chunks to which the file identifiers correspond; and

means for storing an operation log that includes a record of changes to at least one of the namespace data or the mapping data.

14. (canceled)

15. (previously presented) A file system, comprising:

a plurality of servers configured to store files as chunks; and

a master connected to the servers and configured to:

store namespace data that includes file identifiers for which the files are stored as chunks,

store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond,

store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and

store location data that identifies which of the servers stores which of the chunks, where the master is configured to:

determine location information by communicating with the servers, the

location information being based on which of the servers store ones of the chunks,

store the location information as the location data, and
update the location data by periodically communicating with the servers to
obtain changes to the location data.

16. (previously presented) A file system, comprising:

a plurality of servers configured to store file data as chunks; and

a master connected to the servers and configured to:

store namespace data that includes file identifiers for files for which the file data
is stored as chunks,

store mapping data that maps the file identifiers to the chunks to which the file
identifiers correspond,

store an operation log that includes a record of changes to the namespace data and
the mapping data, and

store location data that identifies which of the servers stores which of the chunks,
where the master is configured to:

communicate with the servers to determine location information of the
data, the location information being based on which of the servers store the
chunks, and

store the location information as the location data.

17. (canceled)

18. (canceled)
19. (previously presented) The file system of claim 1, where the file identifiers are organized hierarchically in a tree of directories.
20. (previously presented) The file system of claim 1, where the master stores the namespace data using prefix-compression.
21. (previously presented) The file system of claim 1, where the master is configured to identify one of the chunks via a chunk handle that uniquely identifies the one of the chunks.
22. (previously presented) The file system of claim 21, where the chunk handle encodes a timestamp.
23. (previously presented) The file system of claim 1, where the master is configured to update the location data by periodically instructing the servers to provide information regarding the chunks stored by the servers.
24. (previously presented) The file system of claim 1, where the operation log includes a logical timeline that defines an order for concurrent operations.
25. (previously presented) The file system of claim 1, where the master is configured

to:

determine when a size of the operation log exceeds a threshold, and
create a checkpoint of the operation log when the size of the operation log exceeds the
threshold.

26. (previously presented) The file system of claim 25, where the master is
configured to:

create a new operation log file, and
create the checkpoint as a background operation.

27. (previously presented) The master of claim 13, where the operation log includes
a logical timeline that defines an order for concurrent operations.

28. (previously presented) The master of claim 13, further comprising:
means for determining when a size of the operation log exceeds a threshold; and
means for creating a checkpoint of the operation log when the size of the operation log
exceeds the threshold.

29. (previously presented) The master of claim 28, where the means for creating the
checkpoint includes:

means for creating a new operation log file, and
means for creating the checkpoint as a background operation.

30. (previously presented) A method performed by a master device in a file system that includes the master device connected to a plurality of server devices, the method comprising:

communicating with the server devices to identify file data stored by the server devices as chunks;

storing location information that identifies ones of the server devices that store the chunks;

storing namespace data that includes file identifiers for files for which the file data is stored as chunks by the server devices;

storing mapping data that maps the file identifiers to the chunks to which the file identifiers correspond; and

maintaining an operation log that includes a record of changes to the namespace data and the mapping data.

31. (previously presented) The method of claim 30, where maintaining the operation log includes storing a logical timeline that defines an order for operations including concurrent operations.

32. (previously presented) The method of claim 30, further comprising:

determining when a size of the operation log exceeds a threshold; and

creating a checkpoint of the operation log when the size of the operation log exceeds the

threshold.

33. (previously presented) The method of claim 32, where creating the checkpoint includes:

creating a new operation log file, and
creating the checkpoint as a background operation.

34. (new) The file system of claim 15, where the master is further configured to:
identify one or more of the servers to store a new chunk based on prior chunk
distribution involving the servers, and
place the new chunk at the identified one or more servers.

35. (new) The file system of claim 15, where the master is further configured to:
identify one or more of the servers to store a new chunk based on failure correlation
properties associated with the servers, and
place the new chunk at the identified one or more servers.

36. (new) The file system of claim 16, where the master is further configured to:
identify one or more of the servers to store a new chunk based on prior chunk
distribution involving the servers, and
place the new chunk at the identified one or more servers.

37. (new) The file system of claim 16, where the master is further configured to:
identify one or more of the servers to store a new chunk based on failure correlation
properties associated with the servers, and
place the new chunk at the identified one or more servers.

38. (new) The master of claim 13, further comprising:
means for identifying one or more of the servers to store a new chunk based on utilization
of the servers, prior chunk distribution involving the servers, and failure correlation properties
associated with the servers; and
means for placing the new chunk at the identified one or more servers.

39. (new) The method of claim 30, further comprising:
identifying one or more of the server devices to store a new chunk based on utilization of
the server devices, prior chunk distribution involving the server devices, and failure correlation
properties associated with the server devices; and
placing the new chunk at the identified one or more server devices.